

Atty. Ref.: **42P18469**
Express Mail No.: EV323393136US

UNITED STATES PATENT APPLICATION
FOR

ON-CHIP
FREQUENCY DEGRADATION COMPENSATION

Inventors:
Ravisangar Muniandy
Gregory F. Taylor
Payman Aminzadeh

Prepared by:

BLAKELY SOKOLOFF TAYLOR & ZAFMAN LLP
12400 Wilshire Boulevard, Seventh Floor
Los Angeles, CA 90025-1026
(714) 557-3800

**ON-CHIP
FREQUENCY DEGRADATION COMPENSATION**

FIELD OF THE INVENTION

5 Embodiments of the invention relate generally to on chip testing for integrated circuit reliability and quality assurance, and specifically to on-chip clock frequency compensation for transistor degradation.

BACKGROUND OF THE INVENTION

10 Over time, a semiconductor integrated circuit may become aged and experience higher resistances and lower drive capabilities in its transistors. These effects have become more pronounced as semiconductor integrated circuits have been scaled down.

15 To compensate for the aging effects, a frequency guard band has been used during testing of a new integrated circuit that was recently manufactured. That is, the integrated circuit is tested to operate with a clock rate at a percentage guard band (GB) above the rated frequency. The percentage GB
20 is predetermined through experiments performed during process development. In this manner if the integrated circuit ages accordingly, it has a higher operating probability over time.

 However, not all integrated circuits age in accordance with experiments performed during process development. There
25 is still a small probability that new integrated circuits, with a guard banded frequency rating, will fail as they age. That is, these circuits may fail because they experience greater than expected aging.

 A typical consumer integrated circuit may last seven
30 years, at which point, one percent of their cumulative totals may have experienced catastrophic failure due to the

unexpected degradation of their maximum operating frequencies, F_{max} .

Moreover, some new integrated circuits are robust and can be operated at frequencies greater than the rated operating
5 frequency without ever failing over time. In this case, guard banding the frequency of the integrated circuit to a rated operating frequency is inefficient, as lower than expected aging may occur. These new integrated circuits are never utilized to the extent of their capabilities. The experiments
10 performed during process development to determine the percentage GB for guard banding do not actually reflect how an integrated circuit is used in a system.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 illustrates a block diagram of an integrated circuit including an integrated on-chip frequency compensation block as one embodiment of the invention.

5 Figure 2 illustrates a flow chart diagram of the process to provide frequency degradation measurement and compensation as one embodiment of the invention.

10 Figure 3 illustrates a functional block diagram of the integrated on-chip frequency degradation measurement circuit to provide on chip frequency degradation compensation.

Figure 4 illustrates a functional block diagram of an exemplary embodiment of a reliability oscillator.

Figure 5 illustrates a waveform diagram of waveforms of signals and circuit nodes of the block diagram of Figure 3.

15 Figure 6 illustrates a block diagram of a typical computer system in which embodiments of the invention may be utilized.

20 Figure 7 illustrates a block diagram of a central processing unit in which embodiments of the invention may be utilized.

DETAILED DESCRIPTION

In the following detailed description of embodiments of the invention, numerous specific details are set forth in order to provide a thorough understanding of the invention.

5 However, it will be obvious to one skilled in the art that the embodiments of the invention may be practiced without these specific details. In other instances well known methods, procedures, components, and circuits have not been described in detail so as not to unnecessarily obscure aspects of the
10 embodiments of the invention.

A compensation circuit adjusts the core clock frequency of an integrated circuit to account for real time frequency degradation due to transistor drive current reduction as time passes. By adjusting the core clock frequency of the
15 integrated circuit, product failure may be avoided. The transistor driver current reduces over time as an integrated circuit is used due to known phenomenon caused by p-channel metal oxide semiconductor (PMOS) and n-channel metal oxide semiconductor (NMOS) charge trapping.

20 Data comparisons between actual product frequency degradation closely correlate to frequency degradation in a ring oscillator. The compensation circuit measures the frequency degradation by comparing frequency performance of a DC biased ring oscillator, an AC biased ring oscillator, and
25 an unbiased ring oscillator within the integrated circuit die. The compensation circuit may take these measurements and perform the comparison when the integrated circuit is first powered up or periodically after a time period has passed. If the measurements and comparisons indicate frequency
30 degradation greater than a frequency guard band, the core clock frequency is adjusted. In one embodiment, a new clock ratio (core clock to front-side bus ratio) is provided to compensate for the degradation. In this manner, the

compensation circuit compensates for actual in-use Fmax degradation, improves the overall reliability of the integrated circuit, and provides an opportunity to reduce the frequency guard-bands so that more parts may be sold to
5 operate at higher frequencies.

In one embodiment, an integrated circuit has one or more functional blocks to perform one or more functions and an on-chip frequency compensation circuit. The on-chip frequency compensation circuit includes a selectively enabled
10 reliability oscillator to generate a reference oscillating signal, a clocked reliability oscillator to generate an AC degraded oscillating signal, a static reliability oscillator to generate a DC bias degraded oscillating signal, and a compare circuit coupled to the reliability oscillators to
15 compare the oscillating signals and generate a frequency compensation signal in response to the comparison being greater than a predetermined threshold.

In another embodiment, the on-chip frequency compensation circuit includes a first reliability oscillator, a second
20 reliability oscillator, a third reliability oscillator, a first counter, a second counter, and a third counter. The first reliability oscillator has a selectively powered on ring oscillator to avoid transistor degradation and generates a reference oscillating signal on a first oscillation output.
25 The second reliability oscillator includes a constantly powered clocked ring oscillator to experience AC dynamic transistor degradation and generates an AC degraded oscillating signal on a second oscillation output. The third
30 reliability oscillator includes a constantly powered static ring oscillator to experience DC static transistor degradation and generates a DC bias degraded oscillating signal on a third oscillation output. The first counter has an input coupled to the first oscillation output and generates a reference count

on a first count output. The second counter has an input coupled to the second oscillation output and generates a dynamic count on a second count output. The third counter has an input coupled to the third oscillation output and generates
5 a static count on a third count output.

In another embodiment a method in an integrated circuit is disclosed. The method includes enabling measurement of ring oscillator frequencies of a trio of ring oscillators, measuring a first frequency of a first ring oscillator having
10 dynamically stressed transistors, measuring a second frequency of a second ring oscillator having statically stressed transistors, measuring a third frequency of a third ring oscillator having non-stressed transistors, and comparing the first frequency with the third frequency and the second
15 frequency with the third frequency to determine a measure of transistor degradation.

In yet another embodiment, a microprocessor integrated circuit includes an execution unit to execute instructions; and an integrated on-chip frequency compensation circuit. The
20 integrated on-chip frequency compensation circuit has a reference reliability oscillator to selectively generate a reference oscillating signal, a dynamic reliability oscillator to selectively generate an AC degraded oscillating signal, a static reliability oscillator to selectively generate a DC
25 degraded oscillating signal, a measurement and comparison circuit coupled to the reliability oscillators, the measurement and comparison circuit to receive the reference oscillating signal and the AC degraded oscillating signal to generate a first measure of transistor degradation, the
30 measurement and comparison circuit to receive the reference oscillating signal and the DC degraded oscillating signal to generate a second measure of transistor degradation.

Referring now to Figure 1, a monolithic integrated circuit chip 100 is illustrated as one embodiment of the invention. The integrated circuit 100 includes on-chip frequency compensation circuitry 102, a clock generator 103, and one or more functional circuits 104. The on-chip frequency compensation circuitry 102 is integrated on-chip as part of a monolithic substrate and includes reliability oscillators to perform a characterization of transistor degradation. The on-chip frequency compensation circuitry 102 couples to the clock generator 103. The clock generator 103 couples a clock signal 110 which is coupled to the one or more functional circuits 104. As discussed further below, the integrated circuit 100 may include one or more functional blocks, such as an execution unit for example, and in which case the integrated circuit may be a microprocessor integrated circuit.

The on-chip frequency compensation circuitry 102 may provide a signal to the clock generator 103 to modify an oscillating frequency of a frequency synthesizer or a division ratio of a clock divider. In any case, the on-chip frequency compensation circuitry 102 may alter the frequency of the clock signal generated by the clock generator 103 in order to compensate for transistor degradation and thereby improve the reliability and avoid failure of the integrated circuit over time.

Referring now to Figure 2, a flow chart of the process to provide frequency degradation measurement and compensation is illustrated.

At block 202, the process begins such as upon power-up of the integrated circuit and periodically, after the integrated circuit has been powered-up for a predetermined period of time. For example consider that the integrated circuit is a microprocessor, the degradation measurement and frequency

compensation is done every time the microprocessor is powered-up and can be done periodically (e.g. once a day) thereafter.

At block 204, ring oscillator frequency measurement is enabled. A trio of ring oscillator circuits may be used to
5 determine frequency degradation. Figure 3 illustrates reliability oscillators 320-333 each including a ring oscillator. Two of the trio of ring oscillator circuits, referred to as ring oscillator A and ring oscillator B, are placed under stress while the integrated circuit is powered
10 up. Ring oscillator A receives a DC bias stress without any AC bias stress. That is, ring oscillator A is powered-up in a static state when the integrated circuit is powered up in order to experience a DC bias stress. Ring oscillator B receives an AC bias stress. That is, ring oscillator B is
15 powered-up and is toggled from one logic state to another when the integrated circuit is powered up in order to experience a AC bias stress. The third of the trio of ring oscillator circuits, referred to as ring oscillator C, is used as a reference ring oscillator. Ring oscillator C, the reference
20 ring oscillator, avoids both of the AC bias stress and the DC bias stress. That is, ring oscillator C is only powered-up when the frequency measurement is enabled. This avoids the AC bias stress and the DC bias stress from being applied to ring oscillator C, the reference ring oscillator.

25 At block 206, ring oscillator A's output frequency ("A") is measured and ring oscillator B's output frequency ("B") is measured. Because it receives DC bias stress only, ring oscillator A measures the frequency degradation due to DC bias ("DC bias frequency degradation"). Because it receives AC
30 bias stress, ring oscillator B measures the frequency degradation due to AC bias ("AC bias frequency degradation"). Known methods of measuring the frequency of a ring oscillator may be used to measure ring oscillator A's output frequency

("A") and ring oscillator B's output frequency ("B"), such as that exemplified in U.S. Pat. No. 6,535,013 by Samie B. Samaan, filed December 28, 2000.

At block 208, ring oscillator C's output frequency ("C")
5 is measured. Ring oscillator C, the reference ring oscillator, is unbiased when the integrated circuit is powered-up but for the short time when this measurement is taken. Thus, ring oscillator C measures no frequency degradation from AC bias or DC bias. That is, the oscillating
10 frequency output from the reference ring oscillator, ring oscillator C, is without frequency degradation and exemplifies a fresh circuit when the integrated circuit was first manufactured. Known methods of measuring the frequency of a ring oscillator may be used to measure ring oscillator C's
15 output frequency ("C"), such as that exemplified in U.S. Pat. No. 6,535,013 by Samie B. Samaan, filed December 28, 2000.

Frequency degradation is evident when the measured frequencies ("A", "B") of either ring oscillator A or ring oscillator B is less than the measured frequency ("C") of ring
20 oscillator C.

At block 210, AC bias frequency degradation is compared with DC bias frequency degradation. The greater frequency degradation may then be used to determine a new clock ratio CR_N .

25 A ratio (A/C) of ring oscillator A's output frequency ("A") to ring oscillator C's output frequency ("C") is determined. A ratio (B/C) of ring oscillator B's output frequency ("B") to ring oscillator C's output frequency ("C") is determined. Each ratio is multiplied together with an
30 initial clock ratio CR_0 . The initial clock ratio CR_0 is a clock ratio that is determined for the maximum operating frequency when the integrated circuit is fresh as it was recently manufactured. Then, correlation multipliers, X_1 , X_2 ,

are respectively multiplied to each value to account for a slope difference between frequency degradation experienced by the ring oscillators and the integrated circuit product. The final equation for ring oscillator A is $\{(A/C) * X_1 * CR_0\}$. The
5 final equation for ring oscillator B is $\{(A/B) * X_2 * CR_0\}$. Then the minimum of these two values may be used as the new clock ratio CR_N . These computations and comparisons may be determined by an execution unit within a microprocessor.

At block 212, the new clock ratio CR_N is compared with a
10 guard-banded initial clock ratio ($G * CR_0$) where G is the guard band percentage. For example, a 5% guard band G would be 0.95 (1 minus 5%). If the new clock ratio CR_N is within the guard-banded initial clock ratio ($G * CR_0$) no change may occur. That is the actual clock ratio used by the clock generator 103 to
15 generate a clock signal can remain the same. If the new clock ratio CR_N is outside of the guard-banded initial clock ratio ($G * CR_0$), then the new clock ratio CR_N is output and used as the actual clock ratio by the clock generator 103 to generate the clock signal distributed to the one or more functional blocks
20 104. This comparison may also be determined by an execution unit within a microprocessor.

At block 214, ring oscillator frequency measurement is disabled. Ring oscillator A continues to receive a DC bias stress without any AC bias stress. Ring oscillator B switches
25 at the clock frequency, matching the AC switching of the fastest nodes on the integrated circuit, and thus receives an AC bias stress. Ring oscillator C, the reference ring oscillator, is shut down by being disabled and powered-down so that it avoids both of the AC bias stress and the DC bias
30 stress. In this manner, ring oscillators A and B experience frequency degradation similar to that of the functional blocks of the integrated circuit while ring oscillator C, the reference ring oscillator, is protected from degradation.

At block 216, the process ends until it begins again at block 202 when the integrated circuit is powered-up or a period of time has passed.

Referring now to Figure 3, an embodiment of a frequency degradation measurement circuit 300 is illustrated. The frequency degradation measurement circuit 300 is integrated on-chip and may be included as a part of the on-chip compensation circuitry 102 previously described.

The frequency degradation measurement circuit 300 includes a trio of reliability oscillators 320-322, an inverter 324, an AND gate 326, a trio of prescalers (i.e., dividers) 330-332, a trio of counters 340-342, synchronizers 344-345, AND gate 346, and a compare circuit 350 coupled together as shown and illustrated in Figure 3. A start signal 302, a clock signal 304, and a reset signal 306 are coupled into the frequency degradation measurement circuit 300. A reference count signal 308, a dynamic count signal 310, and a static count signal 312 are coupled into the compare circuit 350. In response to a comparison of the reference count signal 308, the dynamic count signal 310, and the static count signal 312, the compare circuit 350 may generate a frequency compensation signal 355. The frequency compensation signal 355 may be coupled into the clock generator 103 in order that the frequency of the clock signal 110 may be modified.

The circuits of the reliability oscillators 320-322 are substantially similar to each other. How the reliability oscillators 320-322 are connected in the frequency degradation measurement circuit 300 differs.

Note that while the flow chart of Figure 2 shows that the measurement of reliability oscillators A and B is sequential with the measurement of oscillator C in blocks 206 and 208, all three oscillators may be measured at the same time by the

circuit in Figure 3. That is, the order of measurement of the ring oscillators is not critical.

Referring now to Figure 4, an exemplary reliability oscillator 400 is illustrated. The reliability oscillator 400 includes an odd number of inverters connected in series to provide a ring oscillator. In this example, the reliability oscillator 400 includes nine inverters 416A-416I coupled in a series chain together. The reliability oscillator 400 further includes a NAND gate 412, a multiplexer 414, a p-channel field effect transistor (PFET) 418, an n-channel field effect transistor (NFET) 420, and another NFET 422 coupled together as shown and illustrated in Figure 4.

The multiplexer 414 has a pair of inputs, one of which is coupled to an output 405 from the NAND gate 412 and another one coupled to an oscillation output (OSC Out) 410 from the last inverter 416I in the series. The output 415 of the multiplexer 414 couples to the input of the first inverter 416A in the series chain. The multiplexer 414 selects between an output 405 from the NAND gate 412 and an oscillation output (OSC Out) 410 from the last inverter 416I in the series in response to a measurement enable (ME) signal 406. When the ME signal 406 is logically high and power is enabled, the reliability oscillator 400 functions as a ring oscillator with the output of the last inverter 416I coupled into the input of the first inverter 416A in the odd series of inverters. When the ME signal 406 is logically low and power is enabled, the inverter chain in the reliability oscillator 400 may be toggled by the clock input signal 402 if enabled by the enable signal 404 being high. If the enable signal 404 is low, the output 405 of the NAND gate 412 is pulled logically high and no toggling occurs in response to the clock signal 402.

The odd number of inverters in the inverter chain (e.g., inverters 416A-416I) is used as a ring oscillator to measure

device degradation. Instead of inverters, other types of logic gates may be used in the ring oscillator so that they are sized consistently with typical paths on the integrated circuit chip desirable to be monitored, as is the ring
5 oscillator length. Generally, smaller devices are more subject to within die variation as are short chains of gates.

To avoid degradation in the transistors of the ring oscillator, the reliability oscillator 400 has power from the positive power supply or VDD switched into the chain of
10 inverters 416A-416I and the multiplexer 414 in one embodiment. If power# 408, an active low signal, is high, PFET 418 is open and NFETs 420 and 422 are closed so that circuit node 409 is grounded through NFET 420 and the OSC Out 410 is grounded through NFET 422 to the negative power supply or VSS. In this
15 manner, the transistors forming the chain of inverters 416A-416I avoid experiencing AC bias or DC bias degradation. If power# 408, an active low signal, is low, PFET 418 is closed and NFETs 420 and 422 are open so that circuit node 409 is coupled to the positive power supply VDD through PFET 418.
20 With NFET 422 open, the OSC Out 410 may oscillate or be toggled.

In another embodiment, ground from the negative power supply or VSS may be switched into the chain of inverters 416A-416I and the multiplexer 414. In yet another embodiment,
25 both ground from the negative power supply or VSS and power from the positive power supply or VDD may both be switched into the chain of inverters 416A-416I and the multiplexer 414. In either case when power is switched off, the transistors forming the chain of inverters 416A-416I avoid experiencing AC
30 bias or DC bias degradation while other transistors in the integrated circuit are aged and experience degradation.

Logically, the exemplary reliability oscillator 400 can have four logical states with the inverter chain used in one of four ways.

5 State 1: With Power# being logically high, the logical states of the ME input and the enable input are don't cares. The ring oscillator is powered down so that there are no stresses (AC or DC) on the devices that are being measured.

10 State 2: Power# is logically low, ME is logically high, and the logical state of enable is don't care. The ring oscillator can be run so that its period can be measured.

State 3: Power# is logically low, ME is logically low, and the logical state of enable is low. In this state the oscillator is stopped, but a DC bias is maintained across the devices that are to be measured.

15 State 4: Power# is logically low, ME is logically low, and the logical state of enable is high. In this state the microprocessor clock may drive the inverter chain of inverters 416A-416I, subjecting it to the maximum AC stress that any circuit on the chip will experience.

20 Most of the time the three ring oscillators in the frequency degradation circuit are operated with one each in states 1, 3, and 4. This keeps the reference oscillator from degrading, while the other two oscillators are subject to maximum DC and AC degradation.

25 Referring back to Figure 3, the reliability oscillators 320-322 are connected in the frequency degradation measurement circuit 300 differently, each of the reliability oscillators 320-322 being an instance of the reliability oscillator 400.

30 Reliability oscillator 320 is a reference oscillator. Reliability oscillator 320 has its power# input coupled to a done signal 314, its clock input coupled to ground, its enable input coupled to ground, and its measurement enable input (ME) coupled to a measure signal 313. The reliability oscillator

320 is selectively powered on by the done signal 314 when a measurement is to be taken of the reference oscillating frequency. With the clock input and the enable input being grounded, the reliability oscillator 320 does not receive AC
5 bias stress from the clock signal 304.

Reliability oscillator 321 is an AC bias stressed oscillator. Reliability oscillator 321 has its power# input coupled to ground, its clock input coupled to a clock signal 304, its enable input coupled to the positive power supply
10 terminal (VDD), and its measurement enable input (ME) coupled to the measure signal 313. With the power# input coupled to ground, the reliability oscillator 321 is always powered on when the integrated circuit is powered on.

Reliability oscillator 322 is a DC bias stressed
15 oscillator. Reliability oscillator 322 has its power# input coupled to ground, its clock input coupled to ground, its enable input coupled to ground, and its measurement enable input (ME) coupled to a measure signal 313. With the power# input coupled to ground, the reliability oscillator 322 is
20 always powered on when the integrated circuit is powered on. However with the clock input and the enable input being grounded, the reliability oscillator 322 does not receive AC bias stress from the clock signal 304.

Operation of the frequency degradation circuit 300 is now
25 described with reference to Figures 3 and 5. Figure 5 illustrates a waveform diagram of waveforms 501-506 of circuit nodes and control signals of Figure 3 versus time. Waveform 501 illustrates a waveform diagram of the reset signal 306. Waveform 502 illustrates a waveform diagram of the start
30 signal 302. Waveform 503 illustrates a waveform diagram of the stop signal 316. Waveform 504 illustrates a waveform diagram of the done signal 314. Waveform 505 illustrates a waveform diagram of a clock signal 304. Waveform 506

illustrates a waveform diagram of exemplary of the counts 308, 310, and 312.

When the ring oscillators within each reliability oscillator 320-322 are to be measured, the counters 340-342 are first reset by the assertion of the reset signal 306 at 510. The reset signal 306 is then de-asserted at 511 and then the start signal 302 is asserted at 512. The measurement signal output 313 of the AND gate 326 goes high so that the measurement enable input is enabled for each reliability oscillator 320-322. The ring oscillator in each reliability oscillator 320-322 may then start oscillating and the counters 340-342 may then count the number of cycles. Before each counter 340-342, each prescaler 330-332 ("divided by N") reduces the frequency, by dividing by a predetermined number N, of the oscillating signal coupled into the input of each counter. This also reduces the likelihood of metastability in the synchronizer ("Sync") blocks. The divided output of each prescaler 330-332 is respectively coupled into the input of the counters 340-342 for counting.

Counter 340, the reference counter, counts the cycles of the reference ring oscillator of the reliability oscillator 320. Counter 341, the AC stress counter, counts the cycles of the AC biased ring oscillator of the reliability oscillator 321. Counter 342, the DC stress counter, counts the cycles of the DC biased ring oscillator of the reliability oscillator 322. If a counter overflows, the counter overflow ("CO") output is asserted.

The counter 340 counts until overflow is reached and the CO output is asserted as the stop signal 316. The reference oscillator of the reliability oscillator 320 may have a greater oscillating frequency than that of the degraded oscillators in the reliability oscillators 321-322. Thus, it

is expected that counter 340 will overflow prior to the overflow of counters 341-342.

When counter 340 overflows, the stop signal 316 is asserted at 513, which is coupled into the synchronizers 344-
5 345. The stop signal also disables counter 340 from counting further. The stop signal 316 is synchronized with the divided count value from each prescaler 331-332 and coupled to the enable input of counters 341-342 to respectively stop their counting. Each synchronized stop signal is also coupled into
10 the inputs of the AND gate 346 to generate a done signal 314 at 514. The done signal 314 is inverted and de-asserts the measure signal 313 so that the oscillations of the ring oscillators are disabled as the measurement is completed. The logic around the reliability oscillators 320-322 generating
15 the counts 308, 310, 312, including the logic generating the done signal 314 and the stop signal 316, is a state machine.

The count value held by each counter 341 and 342 provides a measure of the degradation from that of the reference ring oscillator. The synchronizers 344-345 allow the counters 341
20 and 342 to count for two extra cycles after the counter 340 of the reference ring oscillator has stopped counting. These two cycles may be subtracted before arithmetic is done to determine if degradation has been significant, or it can be used to simplify the arithmetic in generating a degradation
25 signal 355.

For example, if the threshold of frequency degradation is selected to be 2% degradation, then the counters may be given a length of 100. If either oscillator fails to overflow before it is stopped by the reference oscillator, then it must
30 have degraded by at least 2% (two synchronizer cycles divided by one hundred counts), and the core frequency should be reduced.

In any case, the compare circuit 350 compares the reference count 308 with the dynamic count 310 and the static count 312. If the compare circuit 350 detects that degradation has occurred over a predetermined threshold level, it generates the frequency degradation signal 355. In another embodiment, the compare logic used by the compare circuit 350 to generate the frequency degradation signal 355 may be that previously described with reference to Figure 2.

The frequency degradation signal 355 may be coupled to a clock generator in order to reduce the clock frequency of the clock signal coupled to functional blocks in one embodiment. In another embodiment, the frequency degradation signal 355 may be used to adjust a ratio of clock frequency to bus frequency such as in a microprocessor. By operating the microprocessor at a lower frequency, it can continue to operate reliably in spite of the degraded performance of the transistors that have been aged over time and stress.

Referring now to Figure 6, a block diagram of a typical computer 600 in which the embodiments of the invention may be utilized is illustrated. The computer 600 includes a central processing unit (CPU) 601, input/output devices (I/O) 602 such as keyboard, modem, printer, external storage devices and the like and monitoring devices (M) 603 such as a CRT or graphics display. The monitoring devices (M) 603 provide computer information in a human intelligible format such as visual or audio formats.

Referring now to Figure 7, a block diagram of a typical central processing unit 601 in which the embodiments of the invention may be utilized is illustrated. The central processing unit 601 includes a microprocessor 701 including the embodiments of the invention, a disk storage device 703, and a memory 704 for storing program instructions coupled together. Disk storage device 703 may be a floppy disk, zip

disk, DVD disk, hard disk, rewritable optical disk, flash memory or other non-volatile storage device. The microprocessor 701 and the disk storage device 703 can both read and write information into memory 704 over the memory bus 5 704. Memory 704 is typically dynamic random access memory (DRAM) but may be other types of rewritable storage.

The microprocessor 701 is an integrated circuit 100 in one embodiment as illustrated in Figure 1 including the on-chip characterization circuitry or block 102 and the one or 10 more functional circuits or blocks 104. In this case, at least one of the one or more functional circuits or blocks 104 may be an execution unit to execute one or more instructions. The one or more instructions may be from a software program for example.

15 Utilizing the embodiments of the invention, catastrophic failure of a microprocessor over time due to frequency degradation may be avoided. Embodiments of the invention can improve the prediction of mean times to failure so that a reduction in the frequency guard-band can be explored. 20 Embodiments of the invention provide real-time on-chip degradation measurements to obtain degradation values and in response thereto, the clock frequency can be modified for that particular microprocessor over its lifetime. Embodiments of the invention can reduce the probability of failure due to 25 frequency degradation to nearly zero.

While certain exemplary embodiments of the invention have been described and shown in the accompanying drawings, it is to be understood that such embodiments are merely illustrative of and not restrictive on the broad invention, and that this 30 invention not be limited to the specific constructions and arrangements shown and described, since various other modifications may occur to those ordinarily skilled in the art.